

# Υ-DB: A system for data-driven hypothesis management and analytics

Bernardo Gonçalves  
LNCC – National Laboratory  
for Scientific Computing  
Petrópolis, Brazil  
bgonc@lncc.br

Frederico C. Silva  
LNCC – National Laboratory  
for Scientific Computing  
Petrópolis, Brazil  
fredcs@lncc.br

Fabio Porto  
LNCC – National Laboratory  
for Scientific Computing  
Petrópolis, Brazil  
fporto@lncc.br

## ABSTRACT

The vision of Υ-DB introduces deterministic scientific hypotheses as a kind of uncertain and probabilistic data, and opens some key technical challenges for enabling data-driven hypothesis management and analytics. The Υ-DB system addresses those challenges throughout a design-by-synthesis pipeline that defines its architecture. It processes hypotheses from their XML-based extraction to encoding as uncertain and probabilistic U-relational data, and eventually to their conditioning in the presence of observations. In this demo we present a first prototype of the Υ-DB system. We showcase its core innovative features by means of use case scenarios in computational science in which the hypotheses are extracted from a model repository on the web and evaluated (rated/ranked) as probabilistic data.

## Categories and Subject Descriptors

H.2.1 [Information Systems]: Logical Design

## Keywords

Deterministic hypotheses, design by synthesis, U-relations

## 1. INTRODUCTION

As part of the paradigm shift that makes science ever more data-driven, deterministic scientific hypotheses shall be seen as principles or ideas, which are mathematically expressed and then implemented in a program that is run to give their *decisive* form of data. For a description of the research vision of hypothesis management and its significance, we refer the reader to [8]. In this demo we present a first prototype of the Υ-DB system. We introduce its key research challenges and showcase its main innovative features by means of use cases in computational science.

**Target applications.** Our framework is geared towards hypothesis management applications. Examples of structured deterministic hypotheses include tentative mathemat-

ical models in physics, engineering and economical sciences, or conjectured boolean networks in biology and social sciences. These are important reasoning devices, as they are solved to generate predictive data for decision making in both science and business. But the complexity and scale of modern scientific problems require proper data management tools for the predicted data to be analyzed more effectively.

**Probabilistic DBs.** Probabilistic databases (p-DBs) qualify as such tool, as they have evolved into mature technology in the last decade [19]. One of the state-of-the-art probabilistic data models is the U-relational representation system with its probabilistic world-set algebra (p-WSA) implemented in MayBMS [13]. That is an elegant extension of the relational model we have adopted for the management of uncertain and probabilistic data. It is a core feature of Υ-DB the capability to extract a hypothesis specification and encode it into a U-relational DB seamlessly, ensuring consistency and quality w.r.t. the given hypothesis structure, while alleviating the user from the burden of p-DB design. In short, the system flattens tentative deterministic models into U-relations for data-driven management and analytics.

It comprises some key technical challenges that are dealt with elsewhere [9]. The demonstrated system addresses them throughout a design-by-synthesis pipeline that defines its architecture, see Fig. 1. It processes hypotheses from their XML-based extraction to encoding as uncertain and probabilistic U-relational data, and eventually to their conditioning in the presence of observations.

We introduce next (§2) the pipeline in some detail by emphasizing its main technical challenges. Then in §3 we discuss related work and in §4 the system applicability (assumptions, scope, etc). Finally, in §5 we describe the demo.

## 2. DESIGN-BY-SYNTHESIS PIPELINE

We refer to Example 1 for an accessible and straightforward illustration of the technical challenges addressed by the demonstrated system throughout its pipeline architecture.

**EXAMPLE 1.** *We explore three different theoretical models in population dynamics with applications in Ecology, Epidemics, Economics, etc: (1) Malthus' model, (2) the logistic equation and (3) the Lotka-Volterra model. In the demonstration, such equations are extracted from MathML-compliant XML files (cf. §5). For now, consider that the ordinary differential equation notation ' $\dot{x}$ ' is read 'variable  $x$  is a function of time  $t$  given initial condition  $x_0$ '. These three hypotheses are considered as competing explanations*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

XXX XXX

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.




Figure 1: Design-by-synthesis pipeline for processing hypotheses as uncertain and probabilistic data.

for two phenomena, viz., (1) the nationwide US population from 1790 to 1990, and (2) the Lynx population in Hudson’s Bay, Canada, from 1900 to 1920.  $\square$

$$\dot{x} = bx \quad (1)$$

$$\dot{x} = b(1 - x/K)x \quad (2)$$

$$\begin{cases} \dot{x} &= x(b - py) \\ \dot{y} &= y(rx - d) \end{cases} \quad (3)$$

## 2.1 Hypothesis Encoding

**Structural equations.** In the form of mathematical equations, hypotheses symmetrically relate aspects of the studied phenomenon. For computing predictions, however, deterministic hypotheses are used asymmetrically as *functions* [17]. They take a given valuation over input variables (parameters) to produce values of output variables (the predictions). Such an asymmetry establishes functional dependencies that can be discovered in the predictive data.

Formally, given a system  $\mathcal{S}(\mathcal{E}, \mathcal{V})$  of mathematical equations with a set of variables appearing in them, the AI literature has seen the introduction of an asymmetrical, functional relation among variables that establishes a *causal ordering* [16, 6]. The algorithmic inference of the causal ordering embedded in a deterministic system  $\mathcal{S}$  is a non-trivial challenge for transforming it into a *structural equation model* (SEM), viz., a total mapping  $\varphi_t: \mathcal{E} \rightarrow \mathcal{V}$  from the set  $\mathcal{E}$  of equations to the set  $\mathcal{V}$  of variables. That is a key technical challenge we have addressed in the  $\Upsilon$ -DB system to enable the encoding of a hypotheses as data [9]. It lies at the core of an algorithm we have designed to rigorously transform the hypothesis structure (its equations given in MathML format)<sup>1</sup> into a set of fd’s.

**CHALLENGE 1. (Hypothesis encoding).** Given a deterministic hypothesis  $k$  with its structure  $\mathcal{S}_k$ , extract a total causal mapping  $\varphi_t$  over  $\mathcal{S}_k$  and encode  $\varphi_t$  into an fd set  $\Sigma_k$ .

A validity condition for a given mathematical system to be grounded into target phenomena such that falsifiable predictions can be derived from it is that  $|\mathcal{E}| = |\mathcal{V}|$  [9]. In other words, in addition to the theoretical equations, subsidiary equations (e.g.,  $x_0 = 200, b = 10$ ), must also be provided such that we know the hypothesis (theory) can be grounded into some target phenomenon by means of a simulation trial.

**EXAMPLE 1. (continued).** In this example we are given structures  $\mathcal{S}_k(\mathcal{E}_k, \mathcal{V}_k)$  for  $k = 1..3$  as follows.

- $\mathcal{E}_1 = \{f_1(t), f_2(x_0), f_3(b), f_4(x, t, x_0, b)\};$

<sup>1</sup>MathML is W3C’s standard for encoding mathematical information (<http://www.w3.org/Math/>).

PHENOMENON	$\phi$	Description
1		US population from 1790 to 1990.
2		Lynx population in Hudson’s Bay, Canada, from 1900 to 1920.

HYPOTHESIS	$v$	Name
1		Malthus’ model
2		Logistic equation
3		Lotka-Volterra model

Figure 2: Descriptive (textual) data of Example 1.

- $\mathcal{E}_2 = \{f_1(t), f_2(x_0), f_3(K), f_4(b), f_5(x, t, x_0, K, b)\};$
- $\mathcal{E}_3 = \{f_1(t), f_2(x_0), f_3(b), f_4(p), f_5(y_0), f_6(d), f_7(r), f_8(x, t, x_0, b, p, y), f_9(y, t, y_0, d, r, x)\}.$   $\square$

For instance, the structure of hypothesis  $H_3$  given above has  $|\mathcal{E}_3| = |\mathcal{V}_3| = 9$ . The  $\Upsilon$ -DB system then accepts it as valid and encodes it into fd set  $\Sigma_3$  as shown in Fig. 3 (left). Note, in all fd sets  $\Sigma_k$  with  $k = 1..3$  that  $\phi$  and  $v$  stand (resp.) for the id’s of phenomena and hypotheses. This is a core abstraction of  $\Upsilon$ -DB to capture the data-level semantics of mathematical equations and encode it into fd’s [8].

The fd set  $\Sigma_k$  of hypothesis  $k$  is basic input to schema synthesis at two stages of the pipeline, viz., the ETL and U-intro stages (Fig. 1) we introduce next.

## 2.2 Synthesis ‘4C’

**Design by synthesis.** The problem of design by synthesis has long been introduced by Bernstein in purely symbolic terms as follows [2]: given a set  $U$  of attribute symbols and a set  $\Sigma$  of mappings of sets of symbols into symbols (the fd’s), find a collection  $\mathbf{R} = \{R_1, R_2, \dots, R_n\}$  (the relations) of subsets of  $U$  and, for each  $R_i$ , a subset of  $R_i$  (its designated key) satisfying properties: (P1) each  $R_i \in \mathbf{R}$  is in 3NF; (P2)  $\mathbf{R}$  completely characterizes  $\Sigma$ ; and (P3) the cardinality  $|\mathbf{R}|$  is minimal. More generally, the problem of schema design given dependencies considers criteria [1, Ch. 11]:

- P1’. ( $\simeq$  P1). Desirable properties by normal forms;
- P2’. ( $\simeq$  P2). Preservation of dependencies (“meta-data”);
- P3. The cardinality  $|\mathbf{R}|$  is minimal (minimize joins);
- P4. Preservation of data (the lossless join property).

There is a well-known trade-off between P1’ and P2’, since normal forms that ensure less redundant schemas may lose the property of dependency preservation [1]. In fact, P2’ is important to prevent the DB from the so-called update anomalies, as the fd’s in  $\Sigma_k$  are viewed as integrity constraints to their associated relations. Hypothesis management applications [8], however, are OLAP-like and have an ETL procedure characterized by batch-, incremental-only updates and large data volumes. Thus, it is a design principle of  $\Upsilon$ -DB to trade P2’ for P1’, to favor succinctness (as less redundancy as possible) over dependency preservation (viz., we stick to BCNF). We do that in a principled way by reasoning over the fd’s to ensure the causal ordering is preserved [9]. Also, we shall favor P4 as less joins means faster access to data.

Overall, for such a notion of ‘good’ design, the hypothesis causal ordering mapped into fd set  $\Sigma_k$  needs to be further processed in terms of *acyclic reasoning on its reflexive pseudo-transitive closure*. The demonstrated system relies on an original, efficient algorithm we have designed for that

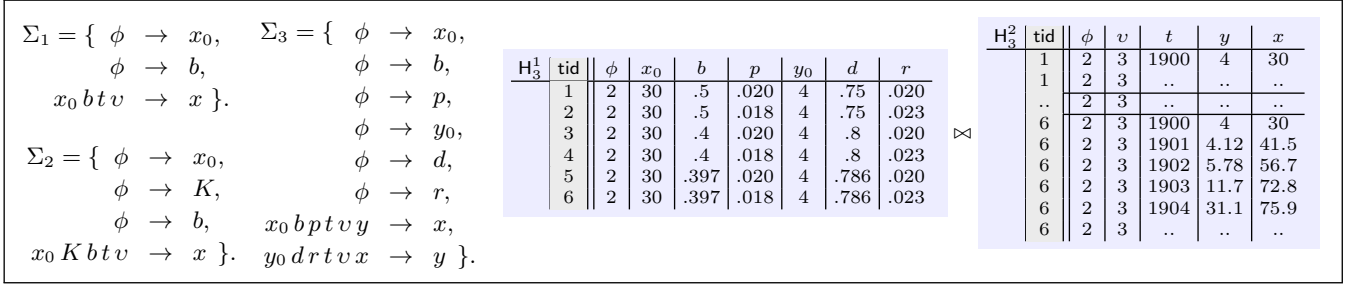


Figure 3: Resources from Example 1. (Left). Primitive fd sets extracted from the given structures  $\mathcal{S}_k(\mathcal{E}_k, \mathcal{V}_k)$  for  $k = 1..3$ . (Right). *Certain* relations  $\mathbf{H}_3 = \{H_3^1, H_3^2\}$  of hypothesis  $v=3$  loaded with trial datasets identified by special attribute  $\text{tid}$ .

[9]. It returns an fd set  $\Sigma'_k$  we motivate and define to be the *folding*  $\Sigma_k^{\rightarrow}$  of  $\Sigma_k$ . Then it applies a variant of Bernstein’s synthesis algorithm (say, for certainty ‘4C’) to render, given  $\Sigma_k^{\rightarrow}$ , a relational schema  $\mathbf{H}_k = \bigcup_{i=1}^n H_k^i$ . Every schema produced this way provably satisfies desirable properties for hypothesis management [9].

**CHALLENGE 2. (Synthesis ‘4C’).** Given fd set  $\Sigma_k$ , derive an fd set  $\Sigma'_k$  (causal ordering processing) to synthesize a relational schema  $\bigcup_{i=1}^n H_k^i$  over it satisfying  $P1'$  (BCNF),  $P3$  and striving for  $P4$  while giving up  $P2'$ .

For hypothesis  $H_3$ , e.g., from  $\Sigma_3$  we derive its folding  $\Sigma_3^{\rightarrow} = \{\phi \rightarrow x_0, \phi \rightarrow b, \phi \rightarrow p, \phi \rightarrow y_0, \phi \rightarrow d, \phi \rightarrow r, \phi t v y \rightarrow x, \phi t v x \rightarrow y\}$ . It is input to  $\mathbf{H}_3 := \text{synthesize}(\Sigma_3^{\rightarrow})$ , which is shown in Fig. 3. Every schema produced by this method provably satisfies the properties just mentioned [9].

Once schema  $\mathbf{H}_k$  is synthesized, datasets  $\mathcal{D}_k^\ell$  computed from the hypothesis under alternative trials (input settings) can be loaded into it to accomplish the ETL phase of the design pipeline. Hypothesis management is then enabled for the user, yet up to the capabilities of a traditional relational DB at this stage of the design pipeline. In Example 1, we consider trial datasets for hypothesis  $v=3$  (viz., the Lotka-Volterra model), which are loaded into the synthesized (certain) schemes in  $\mathbf{H}_3$  (see Fig. 3). Note that the fd’s in  $\Sigma_3$  are violated by relations  $H_3^1, H_3^2$ , but we admit a special attribute ‘trial id’  $\text{tid}$  into their key constraints for a trivial repair (provisionally, yet at the ETL stage of the pipeline) until uncertainty is introduced in a controlled way by synthesis ‘4U’ (U-intro stage, cf. Fig. 1).

## 2.3 Synthesis ‘4U’

There are two sources of uncertainty considered in the  $\Upsilon$ -DB system: (i) *theoretical* uncertainty, arising from the multiplicity of hypotheses targeted at explaining a same phenomenon; and (ii) *empirical* uncertainty, the multiplicity of trials (alternative parameter settings) under a same hypothesis in view of a same phenomenon. It should be expected, then, that the U-intro procedure is operated by the system in the ‘global’ view of all available hypotheses  $k = 1..z$ .

**U-relations.** Synthesis ‘4U’ is performed through a data transformation from ‘certain’ relations to ‘uncertain’ U-relations in the language of probabilistic world-set algebra (p-WSA). The latter comprises the operations of relational algebra, an operation for computing tuple confidence  $\text{conf}$ , and the *repair-key* operation for introducing uncertainty — by giving rise to alternative worlds as maximal-subset repairs of an argument key [13]. For a prompt illustration, consider

query  $Q_0$ , which we present first as p-WSA formula (4), and then in MayBMS’s extended SQL concrete syntax. Note that relation  $H_0$  stores (as foreign keys) all hypotheses available in Example 1 and their target phenomena, see Fig. 4.

$$Y_0 := \pi_{\phi, v}(\text{repair-key}_{\phi @ \text{Conf}}(H_0)). \quad (4)$$

$Q_0$ . **create table**  $Y_0$  **as select**  $\phi, v$  **from**  
(**repair key**  $\phi$  **in**  $H_0$  **weight by**  $\text{Conf}$ );

Query  $Q_0$ ’s result set is materialized into U-relational table  $Y_0$  (Fig. 4). U-relations have in their schema a set of pairs  $(V_i, D_i)$  of *condition columns* [13] to map each discrete random variable  $x_i$  created by the *repair-key* operation to one of its possible values (e.g.,  $x_0 \mapsto 1$ ). The world table  $W$ , inspired in pc-tables [19], stores their marginal probabilities.

The U-intro procedure is based on the fd’s but also on the data. It processes the uncertainty of ‘input’ relations (*u-factorization*) and then propagates it onto ‘output’ relations (*u-propagation*) [9]. The former is motivated by a basic design principle, which is to define exactly one random variable for each actual uncertainty factor (‘u-factor’ for short). The multiplicity of (competing) hypotheses is itself a standard one, the *theoretical* u-factor (cf.  $H_0$  in Fig. 4). The multiplicity of (competing) trials of a hypothesis  $k$  for a phenomenon leads to a problem of learning *empirical* u-factors in the local view of each hypothesis  $\mathbf{H}_k \subseteq \mathbf{H}$ . It is dominated by the problem of fd’s discovery in a relation, which has been extensively studied in the literature (e.g., see [11]).

The system then carries out the same kind of reasoning but over a different fd set, formed by merging the primitive fd set,  $\Sigma_k$ , with the learned (contingent) fd’s [9]. It is input, together with  $\mathbf{H}_k$  itself, to a different synthesis procedure (now, for uncertainty ‘4U’) to render U-relations.

**CHALLENGE 3. (Synthesis ‘4U’).** Given a collection of relations  $\bigcup_{k=1}^z \bigcup_{i=1}^n H_k^i$  loaded with alternative trial datasets  $\bigcup_{k=1}^z \bigcup_{\ell=1}^p \mathcal{D}_k^\ell$ , introduce properly all the uncertainty implicitly present in the database into U-relations  $\bigcup_{k=1}^z \bigcup_{j=1}^m Y_k^j$ .

Fig. 5 shows the rendered U-relations for hypothesis  $v =$

$H_0$	$\phi$	$v$	$Y_0$	$V \mapsto D$	$\phi$	$v$	$W$	$V \mapsto D$	$\text{Pr}$
	1	1		$x_0 \mapsto 1$	1	1		$x_0 \mapsto 1$	.5
	1	2		$x_0 \mapsto 2$	1	2		$x_0 \mapsto 2$	.5
	2	1		$x_1 \mapsto 1$	2	1		$x_1 \mapsto 1$	.33
	2	2		$x_1 \mapsto 2$	2	2		$x_1 \mapsto 2$	.33
	2	3		$x_1 \mapsto 3$	2	3		$x_1 \mapsto 3$	.33

Figure 4: U-relation  $Y_0$  rendered by repair-key in the ‘global’ view for all hypotheses and their target phenomena.

$Y_3^1$	$V \mapsto D$	$\phi$	$x_0$
	$x_2 \mapsto 1$	2	.30

$Y_3^2$	$V \mapsto D$	$\phi$	$b$
	$x_3 \mapsto 1$	2	.5
	$x_3 \mapsto 2$	2	.4
	$x_3 \mapsto 3$	2	.397

$Y_3^3$	$V \mapsto D$	$\phi$	$p$
	$x_4 \mapsto 1$	2	.020
	$x_4 \mapsto 2$	2	.018

$Y_3^4$	$V_1 \mapsto D_1$	$V_2 \mapsto D_2$	$V_3 \mapsto D_3$	$V_4 \mapsto D_4$	$\phi$	$v$	$t$	$y$	$x$
	$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 1$	$x_4 \mapsto 1$	2	3	1900	4	30
	$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 1$	$x_4 \mapsto 1$	2	3	...	...	...
	...	...	...	...	2	3	...	...	...
	$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 3$	$x_4 \mapsto 2$	2	3	1900	4	30
	$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 3$	$x_4 \mapsto 2$	2	3	1901	4.12	41.5
	$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 3$	$x_4 \mapsto 2$	2	3	1902	5.78	56.7
	$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 3$	$x_4 \mapsto 2$	2	3	1903	11.7	72.8
	$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 3$	$x_4 \mapsto 2$	2	3	1904	31.1	75.9
	$x_1 \mapsto 3$	$x_2 \mapsto 1$	$x_3 \mapsto 3$	$x_4 \mapsto 2$	2	3	..	..	..

Figure 5: U-relations rendered for hypothesis  $v = 3$ .

3, whose relations are shown in Fig. 3. Note that  $\text{tid} = 6$  in  $H_3^2$  (Fig. 3) corresponds now to  $\theta = \{x_1 \mapsto 3, x_2 \mapsto 1, x_3 \mapsto 3, x_4 \mapsto 2\}$ , where  $\theta$  defines a particular world whose probability is  $\Pr(\theta) \approx .055$ . This value is computed by the `conf` aggregate operation based on the marginal probabilities stored in world table  $W$ , as illustrated in the demo (cf. §5).

## 2.4 Conditioning

Possible-worlds semantics can be seen as a generalization of *data cleaning*. In the context of p-DB's [19], data cleaning does not have to be (error-prone) one shot [3]. Rather, it can be carried out gradually, viz., by keeping all mutually inconsistent tuples under a probability distribution (ibid.) that can be updated in face of evidence until the probabilities of some tuples eventually tend to zero to be eliminated. This leads us to Remark 1, motivating the hypothesis analytics feature of the demonstrated system.

REMARK 1. Consider U-relational table  $Y_0$  (Fig. 4). Note that it abstracts the goal of a data-driven hypothesis evaluation study as the repair of each  $\phi$  as a key. The  $\Upsilon$ -DB system is designed for enabling users to develop their research with support of query and update capabilities such that they can manage and rate/rank their hypotheses  $v$  w.r.t. each  $\phi$ , until the relationship  $r(\phi, v)$  is eventually repaired to be a function  $f : \Phi \rightarrow \Upsilon$  from each phenomenon  $\phi$  to its best explanation  $v$ .  $\square$

**Bayesian inference.** As suggested by Remark 1, the prior probability distribution assigned via `repair key` is to be eventually conditioned on observed data. The  $\Upsilon$ -DB system enables the user to carry out Bayesian inference steps that update at each step a prior to a posterior. Our primary use case is in computational science (cf. §5), where we have *discrete* random variables mapped to the possible values of (numerical) prediction attributes whose domain are *continuous* (double precision). Therefore Bayesian inference is applied for normal mean with a discrete prior [4].

The procedure uses normal density function (5) with standard deviation  $\sigma$  to compute the likelihood  $f(y|\mu_k)$  for each competing prediction (trial)  $\mu_k$  given observation  $y$ . But in the general case we actually have a sample of independent observed values  $y_1, \dots, y_n$  (e.g., Lynx population in Hudson's Bay over the years). Then, the likelihood  $f(y_1, \dots, y_n|\mu_k)$  for each competing trial  $\mu_k$ , is computed as a product  $\prod_{j=1}^n f(y_j|\mu_{kj})$  of the single likelihoods  $f(y_j|\mu_{kj})$  [4]. Bayes' rule is then settled by (6) to compute the posterior  $p(\mu_k|y_1, \dots, y_n)$  given prior  $p(\mu_k)$ .

$$f(y|\mu_k) = \frac{1}{\sqrt{2\pi}\sigma^2} e^{-\frac{1}{2\sigma^2}(y-\mu_k)^2} \quad (5)$$

Y[s]	$\phi$	$v$	tid	Year	Lynx	Prior	Posterior
	2	1	1	1904	16.49	.167	.047
	2	1	2	1904	18.22	.167	.000
	2	2	1	1904	79.81	.167	.013
	2	2	2	1904	77.82	.167	.017
	2	3	1	1904	89.59	.055	.131
	2	3	2	1904	65.06	.055	.184
	2	3	3	1904	90.08	.055	.124
	2	3	4	1904	77.46	.055	.176
	2	3	5	1904	88.32	.055	.127
	2	3	6	1904	75.92	.055	.180

Figure 6:  $\Upsilon$ -DB query for data-driven hypothesis analytics.

$$p(\mu_k|y_1, \dots, y_n) = \frac{\prod_{j=1}^n f(y_j|\mu_{kj}) p(\mu_k)}{\sum_{i=1}^m \prod_{j=1}^n f(y_j|\mu_{ij}) p(\mu_i)} \quad (6)$$

EXAMPLE 1. (continued). Consider observational data collected from Hudson's Bay from 1900 to 1920 for the Lynx and Hare populations at each year [7]. Suppose that the user queries the hypothesis predictions available in the system (viz., by mapping  $t \mapsto \text{'Year'}$ , and  $x \mapsto \text{'Lynx'}$ ) at, say, year 1904. For rating/ranking the predictions, their (prior) probabilities are conditioned on the available observations and listed (posteriors) for supporting an informed decision about the underlying hypotheses, see Fig. 6. Here, the probabilities are computed taking into account all sample points from year 1900 to 1920, not only the selected year of 1904.  $\square$

The aggregate posterior of each hypothesis as an explanation for a given phenomenon is given by the sum of the posteriors of its alternative predictions, e.g.,  $\Pr_{v \mapsto 2} \approx .030$ .

CHALLENGE 4. (**Conditioning**). Given a prior probability distribution assigned to competing predictions for a target phenomenon and stored as marginal probabilities in the world table  $W$ , apply Bayesian inference and induce effects of posteriors back to table  $W$ .

Note that this is an applied *Bayesian inference* problem that translates into a *p-DB update* one to induce effects of posteriors back to table  $W$ . In this first prototype of  $\Upsilon$ -DB we accomplish it by performing Bayesian inference at application level and then applying p-WSA's `update` (a variant of SQL's `update`) into MayBMS. It is a topic of future work to design a dedicated algebraic operation for conditioning.

## 3. RELATED WORK

**Models and data.** Haas et al. [10] propose a long-term *models-and-data* research program to extend data management technology for 'prescriptive' analytics. This is (sic.) to identify optimal business, policy, investment, and engineering decisions in the face of uncertainty. Such analytics shall rest on deep 'predictive' analytics that go beyond mere statistical forecasting and are imbued with an understanding of the fundamental mechanisms that govern a system's behavior, allowing what-if analyses [10]. They discuss strategies to extend query engines for model execution within a (p-)DB. Along these lines, query optimization is understood as a more general problem with connections to algebraic solvers. Our framework in turn essentially comprises an abstraction of *hypotheses as data* [8]. It can be understood in comparison as putting models strictly into a (flattened) data perspective. Therefore, and as evidenced by this prototype demonstration, it is directly applicable by building upon recent work on p-DBs [19].

**Design by synthesis.** Classical design by synthesis [2] was once criticized due to its too strong ‘uniqueness’ of fd’s assumption, as it reduces the problem of design to symbolic reasoning on fd’s arguably neglecting semantic issues. Hypothesis management, however, sets a use case where design by synthesis is clearly feasible, as it translates seamlessly to data dependencies the reduction made by the user herself into a tentative formal model for the studied phenomenon. In fact, in such a class of applications synthesis methods may be as fruitful for p-DB design as they are for the design of graphical models (cf. [5]).

**Causality in DBs.** Our encoding of equations into fd’s (constraints/correlations) captures the causal chain from exogenous (input) to endogenous (output) tuples. Fd’s are rich, stronger information that can be exploited in reasoning about causality in a DB for the sake of explanation and sensitivity analysis. In comparison with Meliou et. al [15], we are processing causality at schema level. To our knowledge this is the first work to address causal reasoning in the presence of constraints (viz., fd’s).

**Conditioning.** The topic of conditioning a p-DB has been firstly addressed by Koch and Olteanu motivated by data cleaning applications [14]. They have introduced the **assert** operation to implement, as in AI, a kind of knowledge compilation, viz., world elimination in face of constraints (e.g., FDs). For hypothesis management in  $\Upsilon$ -DB, nonetheless, there is a need to apply *Bayes’ conditioning* by asserting observed data (not constraints).

**Hypothesis encoding.** Finally, our framework is comparable with Bioinformatics’ initiatives that address hypothesis encoding into the RDF data model [18]. We point out the Robot Scientist, HyBrow and SWAN (cf. [18]), all of which consist in some ad-hoc RDF encoding of sequence and genome analysis hypotheses under varying levels of structure (viz., from ‘gene G has function A’ statements to free text). Our framework in turn consists in the U-relational encoding of hypotheses from mathematical equations, which is (to our knowledge) the first work on hypothesis relational encoding.

## 4. APPLICABILITY

**Realistic assumptions.** The core assumption of the  $\Upsilon$ -DB system is that the deterministic hypotheses, given in a MathML-compliant XML file, are structured (in equations) such that they are encodable into a SEM that must be valid (i.e.,  $|\mathcal{E}| = |\mathcal{V}|$ ). Also, as a semantic assumption which is standard in scientific modeling, we consider a one-to-one correspondence between real-world entities and variable/attribute symbols within a structure, and that all of them must appear in some of its equations/fd’s. For most science use cases (if not all) involving deterministic models, such assumptions are quite reasonable. It is a topic of future work to explore business use cases as well, and qualitative hypotheses as discussed elsewhere [9].

**Hypothesis learning.** The (user) method for hypothesis generation is irrelevant to our framework, as long as the resulting hypotheses are encodable into a valid SEM. Thus, promising use cases may arise through the incorporation of machine learning methods into our framework to scale up the formation/extraction of hypotheses and evaluate them under the querying capabilities of a p-DB. Consider, e.g., learning the equations from Eureqa.<sup>2</sup>

<sup>2</sup><http://creativemachines.cornell.edu/Eureqa>.

**Model repositories.** Recent initiatives have been fostering large-scale model integration, sharing and reproducibility in the computational sciences.<sup>3,4</sup> They are growing reasonably fast on the web, (i) promoting some MathML-based standard for model specification, but (ii) with limited integrity and lack of support for rating/ranking competing models. For those two reasons, they provide a strong use case for  $\Upsilon$ -DB. The Physiome project, e.g., is planned to integrate very large deterministic models of human physiology [12]. A fairly simple model of the human cardiovascular system has 630+ variables (or equations, as  $|\mathcal{E}| = |\mathcal{V}|$ ), motivating then data-driven hypothesis management and analytics.

## 5. DEMONSTRATION

This first prototype of  $\Upsilon$ -DB is implemented as a Java web application, with the pipeline component in the server side on top of MayBMS (a backend extension of PostgreSQL).

In this demonstration we go through the whole pipeline (Fig. 1), exploring our use case scenarios. In addition to the population dynamics scenario presented throughout this text, the demonstration will include scenarios of computational hemodynamics research, where the hypotheses are extracted from the Physiome model repository, and a scenario extracted from the Virtual Physiological Rat project.<sup>5</sup>

The demonstration will unfold in three phases. In the first phase, we show the ETL process to give a sense of what the user has to do in terms of simple phenomena description, hypothesis naming and file upload to get her phenomena and hypotheses available in the system to be managed as data. In the second phase, we will reproduce some typical queries of hypothesis management (e.g., what hypotheses are available for a selected phenomenon, or listing all the predictions under some selectivity criteria, e.g., predicted values of, say, arterial pressure in systole periods). In the third phase, we enter the hypothesis analytics module. The user will choose a phenomenon for a hypothesis evaluation study, and the system will list all the predictions with their probabilities under some selectivity criteria (e.g., population at year 1920). The predictions are ranked according to their probabilities, which are conditioned on the observational data available for the chosen phenomenon.

Fig. 7 shows screenshots of the system. Fig. 7(a) shows the research projects currently available for a user. Figs. 7(b, c) show the ETL interfaces for phenomenon and hypothesis data definition (by synthesis), and then the insertion of hypothesis trial datasets, i.e., explanations of a hypothesis towards a target phenomenon. Fig. 7(d) shows the interface for basic hypothesis management by listing the predictions of a given simulation trial. Figs. 7(e, f) show two tabs of the hypothesis analytics module, viz., selection of observations and then viewing the corresponding alternative predictions ranked by their conditioned probabilities.

## 6. ACKNOWLEDGMENTS

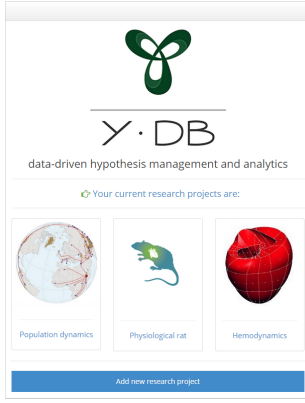
This work has been supported by the Brazilian funding agencies CNPq (grants n° 141838/2011-6, 309494/2012-5) and FAPERJ (grants INCT-MACC E-26/170.030/2008, ‘Nota 10’ E-26/100.286/2013). We thank IBM for a Ph.D. Fellowship 2013-2014.

<sup>3</sup><http://senselab.med.yale.edu/modeldb/>.

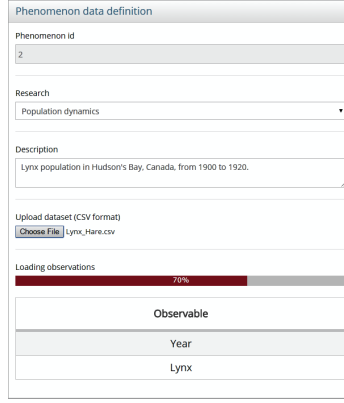
<sup>4</sup><http://physiome.org/>.

<sup>5</sup><http://virtualrat.org/>.

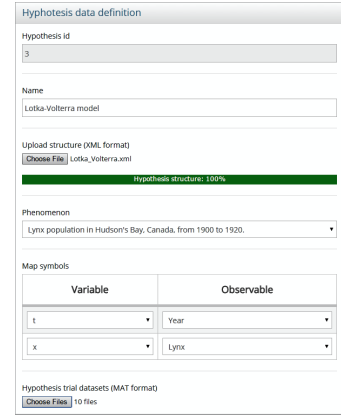




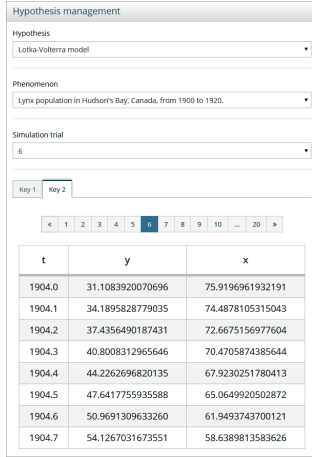
(a) Research dashboard after login.



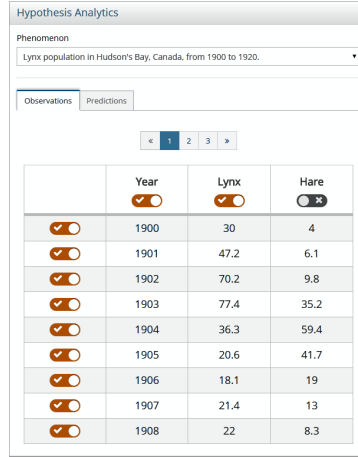
(b) Phenomenon data definition.



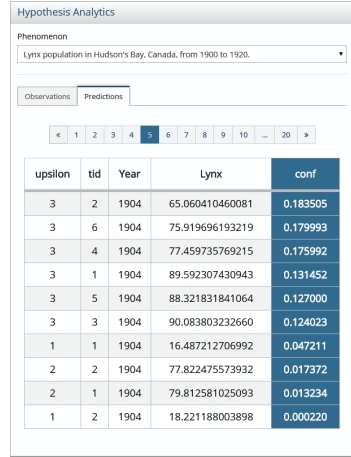
(c) Hypothesis data definition.



(d) Hypothesis management.



(e) Analytics: selected observations tab.



(f) Analytics: ranked predictions tab.

Figure 7: Screenshots of this first prototype of the Y-DB system.

## 7. REFERENCES

- [1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [2] P. Bernstein. Synthesizing third normal form relations from functional dependencies. *ACM Trans. on Database Systems*, 1(4):277–98, 1976.
- [3] G. Beskales, M. A. Soliman, I. F. Ilyas, and S. Ben-David. Modeling and querying possible repairs in duplicate detection. *PVLDB*, 2(1):598–609, 2009.
- [4] W. M. Bolstad. *Introduction to Bayesian Statistics*. Wiley-Interscience, 2nd edition, 2007.
- [5] A. Darwiche. Bayesian networks. *Comm. ACM*, 53(12):80–90, 2010.
- [6] D. Dash and M. J. Druzdzel. A note on the correctness of the causal ordering algorithm. *Artif. Intell.*, 172(15):1800–8, 2008.
- [7] C. Elton and M. Nicholson. The ten-year cycle in numbers of the lynx in Canada. *Journal of Animal Ecology*, 11(2):215–44, 1942.
- [8] B. Gonçalves and F. Porto. Y-DB: Managing scientific hypotheses as uncertain data. *PVLDB*, 7(11):959–62, 2014.
- [9] B. Gonçalves and F. Porto. Design-theoretic encoding of deterministic hypotheses as constraints and correlations in U-relational databases. In *ACM PODS (UNDER REVIEW, CoRR abs/1411.5196)*, 2015.
- [10] P. Haas et al. Data is dead... without what-if models. *PVLDB*, 4(12):1486–9, 2011.
- [11] Y. Huhtala et al. TANE: An efficient algorithm for discovering functional and approximate dependencies. *Comput. J.*, 42(2):100–11, 1999.
- [12] P. J. Hunter and T. K. Borg. Integration from proteins to organs: the Physiome Project. *Nat Rev Mol Cell Biol.*, 4(3):237–43, 2003.
- [13] C. Koch. *MayBMS: A system for managing large uncertain and probabilistic databases*. In C. Aggarwal (ed.), *Managing and Mining Uncertain Data*, Chapter 6. Springer-Verlag, 2009.
- [14] C. Koch and D. Olteanu. Conditioning probabilistic databases. *PVLDB*, 1(1):313–25, 2008.
- [15] A. Meliou et al. Causality in databases. *IEEE Data Eng. Bull.*, 33(3):59–67, 2010.
- [16] H. Simon. *Causal ordering and identifiability*. In Hood & Koopmans (eds.), *Studies in Econometric Methods*, Chapter 3, John Wiley & Sons, 1953.
- [17] H. Simon and N. Rescher. Cause and counterfactual. *Philosophy of Science*, 33(4):323–40, 1966.
- [18] L. Soldatova and A. Rzhetsky. Representation of research hypotheses. *J Biomed Sem*, 2(S2), 2011.
- [19] D. Suciu, D. Olteanu, C. Ré, and C. Koch. *Probabilistic Databases*. Morgan & Claypool Publishers, 2011.